

RESEARCH ARTICLE

Online gendered violence and visibility in electoral autocracies: Evidence from multi-level elections in Türkiye

Jungmin Han¹ , Müge Finkel², Steven Finkel³, Firat Duruşan⁴ and Erdem Yörük⁵

¹Department of Political Science, Trinity College Dublin, Ireland, ²Graduate School of Public and International Affairs, University of Pittsburgh, USA, ³Department of Political Science, University of Pittsburgh, USA, ⁴Center for Computational Social Sciences, Koç University, Türkiye and ⁵Department of Sociology & Center for Computational Social Sciences, Koç University, Türkiye

Corresponding author: Jungmin Han; Email: jhan@tcd.ie

(Received 20 September 2025 UTC; revised 29 May 2026 UTC; accepted 3 June 2026 UTC)

Abstract

Online platforms have become indispensable tools for political candidates to gain visibility, particularly for female candidates in electoral autocracies who often lack access to traditional networks, party hierarchies, and media platforms. We argue that online visibility, while a resource for female candidates, disproportionately exposes them to online violence that functions as a mechanism of political control, reinforcing patriarchal exclusion from the political sphere. Our analysis draws on more than three million tweets directed at 1,834 candidates across the 2023 parliamentary and 2024 local election campaigns in Türkiye, providing a rare opportunity to observe how these dynamics unfold at both the national and subnational levels. We find that while online visibility increases abuse for all candidates, the effect is markedly stronger for female candidates, who are disproportionately targeted across diverse forms of violence, including profanity, demeaning and ridiculing language, physical threats, and ethnic and religious hate speech. Moreover, our results suggest that in local elections, where candidates rely more heavily on personalized campaigns, increases in visibility are more strongly associated with female candidates' exposure to online violence.

Keywords: Online violence; electoral autocracy; gender politics; election; Türkiye

Introduction

Online platforms have become central arenas for electoral politics, enabling candidates to build visibility, mobilize supporters, and bypass traditional gatekeepers (Boulianne 2020; Vaccari and Valeriani 2021). They are especially valuable for female politicians, who often rely on social media to overcome gendered barriers in traditional party organizations and media environments (Osei-Appiah, Mutsvairo, and Orgeret 2023). Yet, these platforms may also reproduce entrenched patriarchal norms and gender stereotypes, exposing female candidates to harassment and abuse from the public (Collignon and Rüdiger 2020; Håkansson 2021; Erikson, Håkansson, and Josefsson 2023). Thus, the very spaces that create opportunities for female politicians may also serve as additional arenas of victimization, paradoxically discouraging their participation.

A growing body of research has begun to map the contours of this paradox, and one finding emerges with consistency: it is not simply being a woman in politics that drives online harassment

but being a visible woman in politics. Across multiple national contexts, studies have found that political prominence – measured by office seniority, media exposure, follower counts, or posting activity – is among the strongest predictors of the volume of online abuse that female politicians receive (Rheault, Rayment, and Musulan 2019; Gorrell, Bakir, Roberts *et al.* 2020; Håkansson 2021; Erikson, Håkansson, and Josefsson 2023). Crucially, this visibility effect is gendered: whereas prominence tends to confer legitimacy and credibility on male politicians, for women it appears to provoke hostility, triggering personal, sexualized, and identity-based attacks at disproportionate rates. Research in social psychology and feminist theory helps explain why women who occupy prominent public roles are often perceived as violating entrenched norms of modesty and deference, and online harassment functions as a mechanism of norm enforcement – a means of punishing those who transgress symbolic gender boundaries (Puwar 2004; Bardall 2013; Krook and Restrepo Sanín 2016). Visibility, in this sense, is not a neutral political resource for female candidates; it is a source of risk.

Yet, this well-established finding has been generated almost exclusively in established democracies. Studies from the United Kingdom (Gorrell, Bakir, Roberts *et al.* 2020), Sweden (Håkansson 2021), and Canada (Rheault, Rayment, and Musulan 2019) have produced important insights, but their institutional contexts, characterized by competitive electoral environments, free media, robust legal protections, and consolidated party systems, differ fundamentally from those prevailing across much of the world. By anchoring the literature in these settings, existing scholarship has implicitly treated consolidated democracy as the baseline condition, leaving online violence in more constrained political environments largely unexamined. This matters especially in electoral autocracies, where weak institutions and patriarchal party structures restrict women's conventional campaigning avenues, making digital platforms among the few accessible spaces for building electoral viability (IFES 2018; Cardo 2021; Chilwa 2022; Matsilele and Nkoala 2023). Here, dependence on digital visibility is not a tactical choice but a structural condition – and if visibility drives gendered harassment, the stakes are considerably higher than existing research has acknowledged. Accordingly, we argue that in contexts of electoral autocracies, online visibility becomes a particular liability for female candidates, as authoritarian movements weaponize digital platforms to target their opponents, with women bearing the heaviest costs of that retribution.

We evaluate this argument using data from the 2023 parliamentary and 2024 local elections in Türkiye, analyzing how candidates' visibility shapes their exposure to online harassment on X (formerly Twitter). As an electoral autocracy marked by shrinking civic space, patriarchal norms, and structural advantages for incumbents (Ugur-Cinar 2023), Türkiye provides a critical case for studying these dynamics. Additionally, comparing parliamentary and local elections allows us to test our argument across contexts that differ in levels of candidate personalization, since parliamentary contests are more party-centered while local elections are more candidate-centered. Our dataset includes more than 3.8 million tweets directed at 1,834 candidates across both campaigns. We measure online offensive and hate speech using a supervised machine-learning pipeline based on the United Nations definition of hate speech (United Nations 2020: 7–8), classifying tweets as offensive or non-offensive and disaggregating by type – physical threats, profanity, demeaning language, and ridicule – as well as by targeted identity such as religion and ethnicity. We operationalize digital visibility through follower counts, which capture the reach of a candidate's audience and exposure to both supporters and detractors (Rheault, Rayment, and Musulan 2019), and control for factors such as candidates' online activities and their incumbency that may relate to both visibility and harassment exposure.

The results show that, while greater visibility in social media increases the likelihood of online abuse for all candidates, its effects are disproportionately severe for women. Across both elections, female candidates experience sharper increases in offensive and hate speech as their follower counts rise. Disaggregated by type of abuse, women are targeted not only with physical threats, ridicule, sexualized insults, and demeaning language, but also with identity-based attacks rooted in religion and ethnicity. These effects are significantly larger in the local elections, where

candidate-centered and socially proximate competition facilitates more individualized forms of online abuse. Moreover, partisan affiliation shapes these dynamics: female candidates from the Green Left Party (YSP/DEM) face substantially higher levels of online violence than their male counterparts or women from other parties, reflecting not only the party's central role as the electoral representative of the Kurdish political movement but also its strong commitment to gender representation, which heightens their visibility and exposes them to intersecting forms of ethnic and gendered abuse.

Beyond extending the analysis from Western democracies to an electoral autocracy, our findings make several contributions. First, we advance scholarship on violence against women in politics by demonstrating that digital visibility functions simultaneously as a resource and a liability. Whereas existing research has largely framed online platforms as empowering for women in politics (Osei-Appiah, Mutsvairo, and Orgeret 2023), we show that it enhances electoral competitiveness while increasing exposure to disproportionate abuse. Second, by operationalizing visibility through follower counts rather than conventional proxies such as position in the political hierarchy (Håkansson 2021), office seniority and incumbency (Gorrell, Bakir, Roberts et al. 2020), or list placement (Holm, Bjarnegård, and Zetterberg 2024), we build on Rheault, Rayment, and Musulan's (2019) approach to provide a more direct and theoretically appropriate measure of online exposure. Our findings systematically replicate the gendered visibility penalty across multiple forms of abuse, corroborating and strengthening one of the literature's most consistent results in a new institutional context. Finally, methodologically, we advance on existing approaches (Rheault, Rayment, and Musulan 2019; Erikson, Håkansson, and Josefsson 2023; Holm, Bjarnegård, and Zetterberg 2024) by moving beyond simple measures of abusive content to differentiate types of offensive and hate speech using supervised machine-learning techniques.

Digital pathways to political participation for women in electoral autocracies

In electoral autocracies, where formal institutions are weakly institutionalized and traditional media often face state control, online platforms provide an alternative sphere for political participation. Digital media lower the costs of engagement, enabling citizens to access information, mobilize, and voice their preferences even when conventional avenues are constrained (Boulianne 2020; Vaccari and Valeriani 2021). Although governments attempt to control online spaces, the accessibility and horizontal structure of social media still expand opportunities for political expression (Gunitsky 2015). In this sense, online platforms function as partial substitutes for restricted offline participation in electoral autocracies, making them crucial sites for political activity.

These participatory benefits extend particularly to women, who face layered barriers to political engagement in less democratic contexts. Patriarchal party structures, limited media representation, and constraints on mobility mean that women are often disregarded in formal political arenas. Online platforms, however, help reduce these disadvantages by providing relatively low-cost, flexible, and less restricted means of communication (IFES 2018; Cardo 2021). For example, during Thailand's 2020 pro-democracy protests, Twitter provided a relatively accessible space for young women activists to articulate grievances and construct collective action frames in a politically restrictive context (Sinpeng 2021). Additionally, South African female politicians use 'metavoicing' and trust-building strategies on social media to enhance their credibility and visibility outside mainstream media (Matsilele and Nkoala 2023). Likewise, women in Nigeria, Ghana, and other electoral contexts leverage online advocacy and generate higher user engagement than men (Yarchi and Samuel-Azran 2018; Chilwa 2022). These cases illustrate how digital campaigning allows women to bypass male-dominated party infrastructures and connect directly with citizens.

Women's reliance on digital platforms in electoral autocracies may stem not only from restricted access to traditional media and male-dominated party structures but also from gendered

expectations that shape how women's political behavior is interpreted. Social Role Congruity Theory (Eagly and Karau 2002) argues that prejudice arises when women occupy roles that are stereotypically associated with masculine traits, creating a perceived mismatch between gender norms and role expectations. In contexts where public authority is coded as masculine and women are expected to remain modest and deferential, female candidates face heightened scrutiny, fewer opportunities within party networks, and limited media coverage. Digital platforms, therefore, become especially valuable for women, providing alternative channels to build visibility and communicate with voters outside institutions shaped by patriarchal norms.

Beyond lowering participation barriers, digital platforms also provide women with both symbolic and substantive forms of representation. Online visibility enables women to occupy public space, articulate policy priorities, and introduce perspectives that are often marginalized in traditional political arenas, thereby broadening the scope of political discourse (Osei-Appiah, Mutsvairo, and Orgeret 2023). Specifically, research shows that women's digital engagement not only increases their descriptive presence but also expands substantive representation by enabling them to raise gendered issues, mobilize supporters, and shape political debates in ways that are less feasible offline (IFES 2018; Yarchi and Samuel-Azran 2018; Cirimele, Macaluso, Agolino *et al.* 2025). Such visibility thus contributes to women's ability to influence agendas and participate more fully in political life.

Online visibility, gender, and the conditional politics of online violence

While scholarship has shown the critical importance of social media for women to engage in effective political participation, a growing body of work increasingly points to the possibility that online visibility also heightens exposure to criticism, harassment, and threats. Research on political violence argues that such forms of abuse can distort core democratic functions by raising the costs of participation, constraining officeholders' behavior, and undermining political representation (Phillips 1995; Krook and Restrepo Sanín 2016; Krook 2017). Yet, existing work leaves open important questions about the conditions under which these vulnerabilities emerge and the actors most likely to bear their consequences.

A growing body of scholarship suggests that the risks associated with online visibility may not be evenly distributed across candidates. Some studies find that women politicians experience higher levels of harassment than comparable men and are subjected to distinct forms of abuse – often sexualized, gendered, or identity-based – that can depress political ambition, harm well-being, and disrupt political engagement (Collignon and Rüdiger 2020; Håkansson 2021; Erikson, Håkansson, and Josefsson 2023). At the same time, empirical evidence on gender differences remains mixed, with some studies reporting little to no systematic gender gap in overall exposure to online harassment (Rheault, Rayment, and Musulan, 2019; Gorrell, Bakir, Roberts *et al.*, 2020; Holm, Bjarnegård, and Zetterberg, 2024). Taken together, these findings suggest that gendered exposure to online violence is likely conditional – shaped by political context, platform dynamics, and moments of heightened visibility – rather than a uniform empirical regularity.

Several scholars further argue that online abuse may reflect deeper resistance to women's political authority, functioning as a mechanism of boundary enforcement that punishes perceived violations of gendered expectations of leadership (Bardall 2013; Krook and Restrepo Sanín 2016). Crucially, among the various factors examined in this literature, one finding emerges with consistency: political visibility is the strongest predictor of the volume and severity of online abuse that female politicians receive. Across multiple studies, women in more prominent roles – whether measured by seniority, media exposure, follower counts, or posting activity – face disproportionately higher levels of harassment than less visible counterparts, and this visibility penalty is systematically steeper for women than for men (Rheault, Rayment, and Musulan 2019; Gorrell, Bakir, Roberts *et al.* 2020; Håkansson 2021; Holm, Bjarnegård, and Zetterberg 2024). These dynamics may become particularly salient during election periods, when candidates face strong

incentives to cultivate visibility to remain competitive and when political competition intensifies public scrutiny and online engagement.

Yet, this well-established visibility finding has been generated almost exclusively in established democracies – contexts characterized by relatively stronger gender equality norms, more independent media environments, and institutional protections that differ in important ways from those found in hybrid or authoritarian regimes (Pacilli and Mannarini 2019; Rheault, Rayment, and Musulan 2019; Gorrell, Bakir, Roberts et al. 2020; Southern and Harmer 2021; Erikson, Håkansson, and Josefsson 2023). This geographic concentration is theoretically consequential: in electoral autocracies, where patriarchal norms are more entrenched, civic space is constrained, and institutional protections for women are comparatively weak, the conditions under which visibility operates as a liability are likely to be more pronounced rather than less. Moreover, relatively few studies explicitly theorize or empirically examine how women's digital visibility during election periods shapes their vulnerability to online abuse, even though elections are precisely when candidates' reliance on digital campaigning is most intense (Gorrell, Bakir, Roberts et al. 2020; Holm, Bjarnegård, and Zetterberg 2024).¹ As a result, we have limited systematic evidence on whether and under what conditions visibility-related penalties emerge during elections in such settings – and whether the most consistent finding in the literature holds or intensifies where democratic protections are weak.

Campaigns, visibility, and online gendered violence in electoral autocracies

To understand why women may face heightened exposure to online violence during election campaigns, it is useful to treat visibility itself as the analytical starting point. Building on Social Role Congruity Theory, which links backlash to perceived misalignments between gendered expectations and role occupancy (Eagly and Karau 2002), political visibility can be understood as a potentially gendered signal of authority and public presence rather than a neutral increase in exposure. Feminist scholarship similarly argues that women who enter highly visible political roles may be evaluated against entrenched norms of modesty, deference, and limited public presence (Puar 2004; Bardall 2013; Krook and Restrepo Sanín 2016). Appearing publicly as a political actor – particularly in the heightened visibility of electoral campaigns – may therefore be interpreted as a symbolic challenge to patriarchal role expectations, creating conditions under which backlash and norm-enforcing forms of online abuse become more likely.

This transgression activates predictable patterns of online violence. Theories of role incongruity and gender backlash suggest that women in positions of authority provoke hostility because their behavior conflicts with stereotypes of femininity, which prescribe warmth and subordination rather than power and leadership (Rudman, Moss-Racusin, Phelan et al. 2012). Whereas men's visibility in politics typically signifies competence and credibility, women's visibility is more likely to be read as illegitimate or threatening (Håkansson 2021). Online abuse thus operates as a social sanction: a way of policing symbolic boundaries, disciplining women who transgress, and reasserting patriarchal hierarchies (Collignon and Rüdiger 2020; Håkansson 2021; Richardson-Self 2021; Erikson, Håkansson, and Josefsson 2023). From this perspective, online violence is an instrument of control that reinforces exclusionary norms about who belongs in politics.

The double standard in how visibility is interpreted becomes especially consequential in electoral autocracies. Elections are moments when visibility is most essential: candidates must be seen and heard to be electorally viable. Yet, women face a paradox. To compete effectively, they must cultivate a strong online presence, but in doing so, they expose themselves to disproportionate abuse. This tension is intensified in electoral autocracies, where elections serve as

¹Holm et al. (2024) provide rare insight into online electoral violence in Tunisia but focuses only on comments on candidates' public Facebook pages and measures visibility indirectly through list placement rather than candidates' actual digital prominence.

instruments of regime legitimation but still generate meaningful incentives for candidates to mobilize supporters (Diamond 2002; Schedler 2002). In such contexts, visibility becomes both indispensable and uniquely precarious. Patriarchal institutions remain entrenched, democratic protections are weak, and online harassment is pervasive and weakly regulated.

Scholars have increasingly conceptualized social media as a political arena that may simultaneously expand opportunities for participation while introducing new forms of risk. From this perspective, digital platforms can enable women to circumvent traditional party hierarchies and mainstream media gatekeepers, potentially allowing them to engage voters more directly and assert political credibility. At the same time, these same platforms may lower the costs of disseminating gendered and misogynistic content, creating new vulnerabilities alongside these opportunities (Theocharis, Barberá, Fazekas *et al.* 2016; Collignon and Rüdiger 2020). Existing studies from a range of contexts suggest that heightened visibility can expose women to disproportionate and often more personalized forms of harassment, though the extent and intensity of these dynamics appear to vary across political settings (Bardall 2013; Rheault, Rayment, and Musulan 2019; Southern and Harmer 2021).

Taken together, these dynamics illustrate our central argument: as women become more visible during elections, they also become more vulnerable to online violence than their male counterparts. Elections heighten the value of digital presence, yet weak institutions and entrenched patriarchal norms magnify its risks. Women cannot compete without cultivating visibility, but cultivating visibility itself exposes them to disproportionate abuse. This paradox lies at the heart of the gendered politics of digital campaigning in weak democracies.

Türkiye as the electoral context for online gendered violence

Electoral autocracy and gender dynamics in Türkiye

Electoral campaigns in Türkiye provide a valuable context for examining the gendered dynamics of online political violence. Türkiye is classified by the Varieties of Democracy (V-DEM) as an electoral autocracy,² where competitive elections are maintained but constrained by shrinking civic space and weakened opposition rights, and one with entrenched patriarchal norms regarding gender roles (Ugur-Cinar 2023, Dalay and Toremark 2024). This combination makes Türkiye an especially illuminating case for studying how political visibility interacts with online harassment under conditions that blend electoral contestation with authoritarian governance.

In Türkiye, these dynamics unfold in a political environment where women's access to political visibility is already structurally constrained within party politics and electoral competition. Women remain persistently underrepresented among elected officials: following the 2023 parliamentary elections, only 118 of 593 members of parliament were women (19.9%) (IPU 2024). This imbalance is evident even at the candidacy stage. Roughly one in four parliamentary candidates in 2023 were women; no woman ran for the presidency, and women received limited visibility during the campaign (OSCE/ODIHR 2023). Similar patterns persist at the local level. After the 2024 local elections, women held only 11 of 81 provincial mayoralities, underscoring the difficulty of translating candidacy into executive authority and sustained public visibility (Reuters 2024).

Beyond formal political institutions, women's public presence in mainstream media is also limited and often shaped through gendered frames. A study of Turkish primetime television news found that 33 percent of news scripts contained no women, and that when women did appear, they were frequently represented as background or instrumental figures or in family-coded roles rather than as authoritative political actors (Erzurum 2013). Against this backdrop, social media platforms provide a comparatively accessible arena for female candidates to generate visibility outside party and media gatekeeping.

²According to V-Dem's Electoral Democracy Index (<http://v-dem.net>), Türkiye has recorded scores below 0.4 since 2015.

Electoral cases: the 2023 parliamentary and 2024 local elections

We focus on Türkiye's two most recent elections: the parliamentary election in May 2023 and the local elections in March 2024. The 2023 parliamentary elections were among the most consequential in Türkiye's recent history. They pitted incumbent President Recep Tayyip Erdoğan and his ruling Justice and Development Party (AKP) against the Republican People's Party (CHP) candidate Kemal Kılıçdaroğlu and a broad opposition coalition. Erdoğan secured re-election in a contested second-round presidential vote, while the AKP-led People's Alliance retained its parliamentary majority. These elections underscored both the persistence of electoral competition and the structural advantages enjoyed by incumbents in hybrid regimes. For female candidates, the high stakes and national visibility of the campaign created intensified exposure to online scrutiny, making the election a critical case for analyzing gendered patterns of harassment.

Additionally, the 2024 local elections provide an important complement to the national contest by shifting the level of competition to subnational politics. Local elections in Türkiye are highly salient, particularly in major metropolitan municipalities such as Istanbul, Ankara, and Izmir, which serve as platforms for opposition parties to challenge the ruling bloc. In 2024, the main opposition party (CHP) achieved a historic victory, winning control of many municipalities such as Bursa and Balıkesir and signaling resilience despite the national-level dominance of the AKP (Silverman 2024).

Analyzing both parliamentary and local elections is central to our theoretical argument. These two electoral arenas differ systematically in the degree to which competition is party-centered versus candidate-centered and in the extent to which candidates are embedded in local social networks. Parliamentary elections in Türkiye are largely structured around party lists and national leaders, with individual candidates given less attention. Local elections, by contrast, are more personalized and socially proximate, heightening the salience of candidates' individual characteristics and public recognition.

These institutional differences generate certain expectations regarding the relationship between candidate visibility and gendered exposure to online violence. In parliamentary elections, the party-centered structure and national framing of competition are likely to dampen the extent to which individual candidates' visibility translates into personalized targeting, limiting the role of gender as a focal point for online abuse. In contrast, local elections place candidates in closer and more continuous interaction with local audiences, increasing personal visibility and activating gendered norms about appropriate political behavior and leadership. As a result, heightened visibility in local elections is more likely to expose female candidates to personalized and norm-enforcing forms of online political violence. We therefore expect the effect of visibility on online violence against women to be larger in local elections than in parliamentary elections.

In contrast to these election-specific dynamics, we expect party affiliation to shape gendered exposure to online violence in a more consistent manner across electoral contexts. Gendered political violence and sexist beliefs are largely shaped by their intersection with other social identities, including race (Hawkesworth 2003; Krook 2017) and partisanship (Cassese and Barnes 2019; Junn and Masuoka 2020). In the Turkish context, candidates from the Green Left Party (YSP, later DEM) – especially women – are likely to face systematically higher levels of hostile online targeting in both parliamentary and local elections. This heightened vulnerability stems from the party's distinctive political positioning along both ethnic and gendered dimensions. As the principal electoral representative of the Kurdish political movement, the YSP/DEM is embedded in a long history of securitization, criminalization, and stigmatization, exposing its women candidates to ethnic and security-oriented forms of online abuse. For example, the arrests and imprisonment of prominent Kurdish women politicians such as Sebahat Tuncel, Gültan Kışanak, Ayla Akat Ata, and Selma Irmak illustrate the broader securitization and criminalization of pro-Kurdish political representation in Türkiye (Dirlik 2016). At the same time, the party's explicit commitment to gender equality, institutionalized

through practices such as co-leadership structures and gender quotas, elevates the political visibility of female candidates and may provoke backlash against women's prominence within a party that openly challenges patriarchal norms.

Political communication through X in Türkiye

Our analysis draws on tweets mentioning candidates' official X handles during the campaign periods for each election. Although existing survey evidence indicates a preference for platforms such as YouTube and Instagram for news consumption among ordinary Turkish citizens (Newman, Fletcher, Robertson *et al.* 2024), X nevertheless occupies a distinctive position as a central arena for political elites, journalists, activists, and politically attentive citizens, giving it an outsized role in shaping the daily political agenda (Sobaci and Karkin 2013). X's critical role in Turkish election campaigns has been documented since the early 2010s. Metropolitan mayoral candidates have made frequent use of the platform for campaign communications (İkiz, Sobaci, Yavuz *et al.* 2014), and research shows that X is the platform on which prominent political leaders in Türkiye have historically accumulated the largest followings relative to other social media platforms (Darı 2018). Platform audience metrics underscore this 'influential minority' character: data from X's advertising resources indicate roughly 19.7 million users in Türkiye in early 2025, suggesting that while X is smaller than mass-usage platforms such as Facebook, it remains a consequential arena for election-period political contention (Kemp 2025).

Our own data are consistent with these findings, confirming that X remains salient for political interactions between politicians and the public. During the parliamentary elections, nearly 70 percent of candidates maintained active X accounts, defined as publicly accessible accounts that had posted at least once during the time of data collection. The total number of posts from these accounts exceeded eight million, with a median of 4,735 posts per account. The number of times these accounts were mentioned in tweets analyzed by our models was approximately 3.74 million, with a median of roughly 300 mentions per account.

It is possible that toxicity on X increased after Elon Musk's 2022 acquisition of the platform, given the subsequent rollback of content moderation and abuse-prevention measures framed as efforts to promote free speech (Hickey, Schmitz, Fessler *et al.* 2023). However, studies documenting these tendencies have largely focused on English-speaking contexts, where the political symbolism surrounding Elon Musk's ownership may carry greater salience than in Türkiye. So, while heightened exposure to offensive content on X warrants careful consideration, the size of X's user base and its prominence in elite political communication renders it an analytically appropriate site for examining the dynamics of online political violence in the Turkish electoral context.

Research design

Data collection

We collected data on online violence against electoral candidates and their account characteristics at two points in time: September 2023, approximately four months after the parliamentary election, and May–June 2024, a few weeks after the local elections. For both elections, the universe of candidates was identified using the official candidate lists published by the Yüksek Seçim Kurulu (YSK), Türkiye's Supreme Election Board (<https://www.ysk.gov.tr/>). In total, these lists yielded 1,984 male and 651 female candidates for the 2023 parliamentary election, and 670 male and 738 female candidates for the 2024 local elections.

Next, we identified candidates' X handles through two independent workflows designed to minimize misidentification. First, we conducted candidate-by-candidate web searches combining each candidate's full name and manually reviewed the most plausible results. Second, an independent research assistant conducted a separate manual search directly through the X interface. The results of the two searches were subsequently cross-checked and reconciled, with

discrepancies resolved through manual inspection. For each retained X account, we verified authenticity using profile- and content-level indicators consistent with an official campaign presence. These indicators included alignment between the account name, username, biographical information, profile imagery, links to official party or campaign pages where available, and campaign-relevant posts referencing the candidate's party affiliation and electoral district. When multiple plausible accounts existed or identification remained ambiguous, we excluded the candidate rather than risk including a same-name user, parody, or impersonation account. Because our design relies on publicly observable interactions, the final sample includes only accounts that were publicly accessible at the time of collection; protected, deleted, suspended, or otherwise unavailable accounts were not included. The final dataset covers tweets directed at 954 parliamentary candidates (553 men and 401 women) and 880 local mayoral candidates (423 men and 457 women).

The time frame of our analysis was selected to reflect campaign dynamics. For local elections, data collection covers the final three months prior to election day, as mayoral candidacies in Türkiye are often formalized relatively late and candidates rarely engage in active campaigning before official nominations, whereas parliamentary candidates frequently begin campaigning earlier, justifying a longer six-month window. To address concerns about campaign duration and temporal consistency, we replicate all analyses using the same three-month window for both elections; the results are reported in the online Appendix 5. Finally, all statistical models include controls for time to account for temporal variation in exposure and ensure that the estimated effects are not driven by differences in the timing of online activity.

Tweets referencing the candidates were gathered at a time when academic access to the platform's official Application Programming Interface (API) had become highly restricted. To address these constraints, recent computational studies have relied on the Nitter web interface (Alzahrani and AlGhamdi 2025; Ortiz, Ahmed, Salas et al. 2025). Nitter is an open-access, privacy-oriented alternative front end that allows researchers to access and extract publicly available content from X without requiring authenticated API access (www.nitter.net). Employing Nitter facilitates the systematic and reproducible collection of public political communication while minimizing exposure to platform-imposed access limitations and algorithmic personalization effects. Because collection occurred after the elections, some offensive or abusive content may have been deleted by users or removed through platform enforcement prior to retrieval. We therefore interpret observed levels of abusive content as conservative lower bounds and discuss the implications of moderation-related missingness in the limitations section below.

The dataset includes a total of 3,656,433 tweets from the 2023 Turkish parliamentary election and 321,394 tweets from the 2024 local elections. For the parliamentary election, we collected tweets directed at 954 parliamentary candidates. These candidates received 3,229,023 tweets mentioning male candidates and 427,410 tweets mentioning female candidates. For the local elections, the dataset covers 880 mayoral candidates, with 132,156 tweets directed at male candidates and 189,238 tweets directed at female candidates. This indicates that, while male candidates attracted most of the online attention in the 2023 parliamentary election, the pattern reversed in the 2024 local elections, where female candidates received more tweets than their male counterparts. These descriptive differences reflect variation in patterns of online attention across electoral contexts rather than a selection mechanism specific to candidate gender.

Categorizing online offensive and hate speech

Following much of the existing literature on online political violence, the core analytical task of this study is to categorize tweets that reference electoral candidates as containing either offensive or hateful language. At a broad level, we understand offensive speech as language intended to insult, demean, or intimidate, and hate speech as language that targets individuals or groups based

on protected identity characteristics. At the same time, we recognize that this task is fraught with significant conceptual and empirical challenges. Definitions of offensive and hateful speech are contested, context-dependent, and often difficult to operationalize in a consistent and transparent manner.

Our study examines the extent to which tweets directed at electoral candidates during the campaign period contain such forms of violent language. Offensive and hateful speech represent one category of online violence against women in politics, alongside other forms such as cyberstalking, sexual harassment, doxing, intimate image abuse, and disinformation (Krook 2020). Language intended to cause offense can function as a means of attacking and intimidating female candidates, while sexist and misogynist tropes may also be deployed as implicit ‘dog whistles,’ enabling harassment while maintaining plausible deniability on the part of perpetrators. Detecting both explicit and implicit forms of abuse requires analytical models that are sensitive to cultural, political, and linguistic context. It also requires a conceptual framework that is sufficiently flexible to capture the multidimensional nature of offensive and hateful speech, yet sufficiently precise to generate categories that are clearly delimited, robust, and practically applicable as document labels. Considering this conceptual diversity, we construct a multi-layered ontology of offensive and hateful speech that encompasses multiple types within each category of violence.

For our overarching definition of online violence, we draw on the working definition provided by the United Nations Strategy and Plan of Action on Hate Speech, which defines hate speech as ‘any kind of communication in speech, writing or behavior, that attacks or uses pejorative or discriminatory language with reference to a person or a group based on who they are, in other words, based on their religion, ethnicity, nationality, race, color, descent, gender or other identity factor’ (United Nations 2020: 7–8). This definition highlights two core dimensions that guide our classification framework: the type of offense (including attacks, discriminatory language, and pejorative expressions) and the identity factors targeted (such as gender, ethnicity, race, or religion).

Our ontology treats these dimensions as modular and analytically distinct, combining them to form comprehensive clusters of offensive and hateful speech. Rather than relying on a singular detection model, we develop separate classification tools for each component of the broader constructs of offensive and hate speech. The following section presents the specific categories used in our analysis, along with illustrative examples drawn from the dataset.³

Offensive speech: *Words or phrases that are intended to assault, threaten, insult, ridicule, or belittle a particular group, individual, or institution.*

- (a) **Physical harm:** *Language that threatens or incites physical violence or serious harm to an individual or group, including calls for persecution or repression that deprives the targets of rights.*
- (b) **Profanity:** *Language that insults individuals or groups through obscenities or expressions considered morally repugnant.*
- (c) **Demeaning:** *Language that demeans, degrades, or belittles a targeted individual or group’s personality, moral values, or humanity. Discriminatory language falls within this category.*
- (d) **Ridicule:** *Language that mocks, ridicules, or shames individuals or groups. While such language may appear harmless in familiar settings, it can function as bullying and cause harm.*

Hate speech: *Any offensive language, as defined above, that targets an individual or social group based on their identity, such as their gender, ethnicity, race, or religion.*

³Additional details on the annotation procedures and examples of the different types of online violence are provided in online Appendix 1.

- (a) **Gender-based hate:** *Offensive speech targeting gender identity, including sexist or discriminatory language such as gender stereotyping.*
- (b) **Religion-based hate:** *Offensive speech targeting individuals or groups based on their beliefs, including disparagement based on religiosity or non-religiosity, alleged heresy, or membership in a faith or denomination.*
- (c) **Ethnicity-based hate:** *Offensive speech targeting individuals or groups based on their ethnic, cultural, or national identity.*

Using these interrelated concepts as an ontology for automatic detection has several advantages. First, it delimits different levels of objectivity by involving categories that depend on explicit intent (such as insult or threat) or less explicit innuendos or ‘dog-whistles’ in categories (demeaning and mocking). Second, the categories can be arranged hierarchically according to the intensity of the offense, from physical threat at the apex to simple mockery at the bottom, enabling more complex analyses such as those linking levels of violence with social group differentiation. Finally, by independently detecting speech targeting gender, religious, and ethnic identities and combining these categories into a comprehensive measure of hate speech, we ensured that no identity group represented in the data was excluded from the definition.

The supervised models were trained and tested on a gold-standard corpus (GSC) containing 4,000 documents sampled from the tweets posted before the relevant election that contained candidate handles. Gold-standard corpora are high-quality training datasets that contain documents annotated by at least two research assistants in accordance with an annotation scheme that is predefined and tested to ensure a high level of accuracy and consistency (Hürriyetoglu, Tanev, Zavarella et al. 2021). In our research design, GSC documents were sampled from the target dataset of tweets that mentioned electoral candidates in a way that normalized party and gender distributions to minimize selection bias.

Because the distribution of candidates is highly imbalanced across political parties and between genders, a purely random draw from the tweet pool would have disproportionately captured highly visible male politicians, resulting in a skewed training set that underrepresents less prominent candidates. To address this concern, we first selected near-equal numbers of candidates across major political parties and gender groups. From the tweets mentioning this pre-selected set of candidates, we then randomly sampled 3,000 tweets to form the initial training dataset. This balanced sample was annotated by three graduate students of social sciences from Koç University based on our annotation manual.⁴ The annotators were trained on the ontology of offensive and hateful speech by the domain expert, who then supervised the annotation process via weekly meetings with the annotators. To assure the gold-standard quality of the training data, each of the tweets in the 4,000-strong set was labeled by two annotators. The instances where the annotators disagreed were adjudicated by the domain expert. The agreed-upon and adjudicated labels became the final labels of the annotation process on whose basis the machine-learning models were trained. On average, inter-annotator agreement was monitored and found to exceed 95 percent on average.

The GSC was then used to train the hate speech and offensive speech classification models by fine-tuning a pretrained transformer model on the Turkish language (Turkish BERT Model).⁵ The macro-F1 scores yielded by the offensive speech and hate speech classification models were 0.70 (precision = 0.73; recall = 0.69) and 0.77 (precision = 0.88; recall = 0.71), respectively. These values are moderate to high compared with state-of-the-art results reported for English-language models (Fortuna and Nunes 2018; Alkomah and Ma 2022), though they could have been improved by increasing the training data size. The model errors were evaluated to see whether models had systematic bias, particularly in the case of hate speech models where there was a discrepancy

⁴See Appendix 1 for more details regarding the online offensive/hate speech annotation manual.

⁵See <https://huggingface.co/dbmdz/bert-base-turkish-128k-cased> for more details about the Turkish BERT Model.

between the recall and precision scores. Based on the error analysis, no systematic bias against candidate groups was observed along political party or gender lines. Descriptively, while the total number of tweets was much larger during the 2023 parliamentary election, the 2024 local elections saw a higher proportion of abusive content. In the parliamentary campaign, 492,329 tweets (15.56%) were coded as offensive and 81,424 (2.28%) as hate speech out of 3,656,433 tweets. By contrast, in the local campaign, 51,386 tweets (19.03%) were classified as offensive and 13,639 (4.43%) as hate speech from a total of 321,394 tweets.

Online visibility of electoral candidates

We use X follower counts as a direct indicator of visibility. Follower counts capture the size of a candidate's potential audience and shape the likelihood that their posts will be circulated, debated, or attacked. In electoral contexts where X functions as a key site of political communication, follower counts reflect both the reach of a candidate's message and the intensity of public attention directed toward them (Rheault, Rayment, and Musulan 2019). This measure is particularly relevant for analyzing online abuse, since a greater number of followers increases exposure not only to supporters but also to critics, trolls, and partisan opponents. Operationalizing visibility in this way embeds the concept within the very platform where harassment occurs, ensuring conceptual and empirical alignment. We admit that follower count may not be a perfect proxy in that it captures potential rather than realized exposure – a large follower base does not guarantee that every follower is active or engaged. While inactive followers may cause follower count to overstate true audience reach, this is unlikely to introduce systematic bias into gender comparisons, as there is no strong reason to expect inactivity to be distributed unevenly across male and female candidates.

Alternative indicators of digital visibility, such as incumbency status or the volume of tweets produced by a candidate's account, could also plausibly capture aspects of public exposure. Incumbency, for instance, is often associated with higher name recognition and media attention. For this reason, our statistical models incorporate incumbency as a control variable to account for this potential source of exposure. However, incumbency has important limitations in the present context. It is an institutionally defined attribute that does not map cleanly onto platform-specific visibility and varies substantially in its relevance across national and local electoral arenas. By contrast, follower counts provide a platform-native measure of visibility that reflects accumulated audience reach and structures the likelihood of both engagement and harassment, making them better suited for analyzing gendered exposure to online political violence.

On the other hand, tweet counts capture a candidate's level of online activity, but they have important limitations as a measure of visibility. Tweet volume conflates exposure with behavior: candidates who post frequently are not necessarily more visible to broader audiences, and high tweet counts may reflect strategic communication styles or resource constraints rather than genuine public attention. Unlike follower counts, which directly index accumulated audience reach, tweet counts say little about how many users can potentially encounter or amplify a candidate's content. That said, tweet volume is not entirely without informational value. As follower counts increase, the reach of each individual tweet expands correspondingly, meaning that posting activity translates into meaningful visibility only conditionally on audience size. We therefore retain tweet volume as a control variable in our main models while treating follower count as the primary measure of platform visibility. To further assess whether high tweet activity distorts our visibility measure, we also conduct supplementary analyses restricted to candidates whose tweet counts fall below the sample mean – a subsample in which follower count more cleanly captures audience reach without the confounding influence of unusually prolific posting behavior. These results are reported in Appendix 6.

In our datasets, the visibility variable – measured by candidates' follower counts on X – is skewed in both elections. In the 2023 parliamentary election, follower counts range from as few as

3 to over 6.9 million, with a mean of 67,301, reflecting a small number of highly visible candidates alongside a majority with modest audiences. For example, only Süleyman Soylu had more than three million followers among parliamentary candidates. A similar pattern emerges in the 2024 local elections, where follower counts range from 0 to 1.15 million, with a substantially lower overall mean of 6,169. Only two mayoral candidates, such as Meral Daniş Beştaş and Erkan Baş, exceeded 500,000 followers. To assess the sensitivity of our findings to this skewness, we replicate the analyses excluding candidates in the upper and lower 5 percent of the follower-count distribution; the results are reported in Appendix 7.

Control variables

As mentioned above, our models include controls for general online engagement and incumbency. Specifically, we account for communication intensity by including the number of tweets posted by the candidate during the campaign period, as well as the number of accounts followed to capture differences in platform use. In the 2023 parliamentary models, we additionally control for incumbency status, given that officeholders may attract distinct forms of attention and criticism.⁶ All variables are measured at the individual level and appropriately rescaled for interpretability.

Temporal closeness to election day represents another crucial factor influencing candidate visibility and online abuse. As campaigns intensify, candidates experience a surge in attention, with social media interactions peaking in the final weeks before voting. Controlling for the timing of tweets relative to election day allows us to separate general patterns of online violence from dynamics specifically tied to the electoral calendar. In this way, our analysis accounts for the cyclical rhythms of campaigning that affect candidates' visibility and vulnerability in the digital sphere. We measure this variable on a daily basis and rescale it in 30-day units to facilitate substantive interpretation in the statistical models.

Finally, we control for party affiliation in our main analyses to account for baseline ideological differences across candidates.⁷ Party labels in Türkiye capture meaningful variation in political orientation and electoral positioning that may shape exposure to online harassment independently of candidate visibility. Following the discussion above, we then examine party-based differences in the effects of visibility on online violence in a separate set of analyses, allowing us to assess how the relationship between visibility and online harassment varies across parties rather than absorbing these differences through controls alone.

Results

We begin our analysis by assessing the baseline relationship between candidate gender and exposure to online violence. Tables 1 and 2 report estimates from logistic regression models predicting offensive and hate speech directed at candidates in Turkish elections. Table 1 presents results from the 2023 parliamentary election (Models 1–4), while Table 2 reports findings from the 2024 local elections (Models 5–8). Across both sets of models, we evaluate the effects of gender and online visibility on the probability of receiving offensive or hate speech, controlling for a range of covariates. The results reveal no uniform gender effect across electoral contexts. Averaged effects indicate that during the 2023 parliamentary election (Models 1 and 3), women were less likely than men to be targeted, with a 6.85 lower percent probability of offensive speech and a 1.06

⁶In the local election models, the incumbency variable was not available because the official candidate rosters do not provide reliable information on candidates' incumbent status (<https://www.ysk.gov.tr/>).

⁷For the local elections, we identified candidates from 33 parties. For analytical simplicity, we categorized parties whose candidates received fewer than 3,000 targeting tweets in total or had fewer than five male or female candidates as "Others" in our dataset.

Table 1. Logistic regression models on online violence in the 2023 parliamentary election

	<i>Offensive speech</i>		<i>Hate speech</i>	
	Model 1	Model 2	Model 3	Model 4
Female	−0.617*** (0.006)	−0.746*** (0.008)	−0.512*** (0.014)	−0.603*** (0.019)
Follower count (100K)	0.008*** (0.0001)	0.008*** (0.0001)	0.004*** (0.0002)	0.004*** (0.0002)
Female × follower count		0.056*** (0.002)		0.036*** (0.005)
Following count (100K)	3.436*** (0.039)	3.393*** (0.039)	5.366*** (0.077)	5.339*** (0.077)
Tweet count (100K)	0.082*** (0.011)	0.107*** (0.011)	−0.617*** (0.022)	−0.602*** (0.022)
Time to election (1 month)	−0.01*** (0.001)	−0.011*** (0.001)	−0.082*** (0.003)	−0.084*** (0.003)
Incumbency	0.196*** (0.004)	0.174*** (0.004)	0.606*** (0.009)	0.591*** (0.010)
(Intercept)	−2.191*** (0.004)	−2.190*** (0.004)	−4.098*** (0.009)	−4.096*** (0.009)
Party fixed effects	0	0	0	0
AIC	2470787	2470206	657876	657825
BIC	2470930	2470362	658019	657981
Log likelihood	−1235382	−1235091	−328927	−328900
Num. obs.	3278887	3278887	3278887	3278887

Note: Logistic regression models for both elections were estimated using glm in R.

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

Table 2. Logistic regression models on online violence in the 2024 local election

	<i>Offensive speech</i>		<i>Hate speech</i>	
	Model 5	Model 6	Model 7	Model 8
Female	0.145*** (0.012)	−0.025 (0.013)	0.058** (0.022)	−0.161*** (0.024)
Follower count (100K)	0.083*** (0.002)	0.039*** (0.003)	0.144*** (0.005)	0.069*** (0.006)
Female × follower count		0.235*** (0.007)		0.282*** (0.013)
Following count (100K)	−0.640*** (0.002)	−0.435** (0.159)	−3.612*** (0.589)	−3.140*** (0.578)
Tweet count (100K)	0.405*** (0.034)	0.065 (0.036)	0.189* (0.074)	−0.358*** (0.081)
Time to election (1 Month)	0.135*** (0.006)	0.123*** (0.006)	0.162*** (0.011)	0.502*** (0.035)
(Intercept)	−1.894*** (0.015)	−1.792*** (0.015)	−2.914*** (0.024)	−2.780*** (0.024)
Party fixed effects	0	0	0	0
AIC	254747	253679	100953	100471
BIC	254949	253892	101155	100683
Log likelihood	−127354	−126819	−50457	−50215
Num. obs.	307433	307433	307433	307433

Note: Logistic regression models for both elections were estimated using glm in R.

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

percent lower probability of hate speech. In contrast, the 2024 local elections (Models 5 and 7) display the opposite pattern, with female candidates facing higher risks of abuse, including increases in the probability of 1.86 percent in offensive speech and 0.15 percent in the probability of hate speech.




Figure 1. Online visibility and gender differences in online violence.
 Note: Error bars in the coefficient plots represent 95% confidence intervals.

These findings help clarify previously mixed evidence in the literature. While prior studies report divergent results regarding gender differences in online harassment (Rheault, Rayment, and Musulan 2019; Collignon and Rüdig 2020; Gorrell, Bakir, Roberts et al. 2020; Håkansson 2021; Erikson, Håkansson, and Josefsson 2023; Holm, Bjarnegård, and Zetterberg 2024; Finkel, Thakur, Finkel et al. 2025), our results show that gendered exposure is contingent on electoral context. In Türkiye's party-centered parliamentary elections, attention is directed primarily toward parties rather than individual candidates, reducing the salience of gender as a basis for targeting. In contrast, the more candidate-centered and locally embedded nature of municipal elections increases personal visibility and activates localized gender norms, thereby heightening women's vulnerability to personalized and norm-enforcing abuse. These patterns support the argument that gendered online violence is conditional rather than uniform, varying systematically across electoral arenas.

Despite these contextual differences, a consistent pattern emerges regarding the role of online visibility. Across both electoral settings, increases in visibility have a significantly stronger effect on women's likelihood of experiencing offensive and hate speech than on men's. We capture these gender-differentiated dynamics using interaction terms between gender and online visibility in the logistic regression models. In the 2023 parliamentary election, the interaction terms in Models 2 and 4 of Table 1 are positive and statistically significant, indicating that visibility amplifies women's risk of receiving both offensive and hate speech to a greater extent than it does for men. The same pattern appears in the 2024 local elections, where the interaction terms in Models 6 and 8 of Table 2 are likewise positive and statistically significant.

Figure 1 illustrates these results by plotting predicted probabilities derived from the regression models. Based on average marginal effects calculated over the observed sample, a one-hundred-thousand-follower increase during the 2023 parliamentary election raises the probability of offensive speech by 0.09 pp (percentage points) for men but by 0.54 pp for women and increases the probability of hate speech by 0.01 pp for men compared to 0.1 pp for women. Evidence from the 2024 local elections shows a similar pattern: the same increase in followers is associated with a 0.47 pp rise in offensive speech for men versus a substantially larger 3.48 pp increase for women and a 0.25 pp increase in hate speech for men compared to 1.66 pp for women. Overall, these findings demonstrate that while online visibility generally intensifies exposure to online violence,




Figure 2. Online visibility and gender differences across offensive speech types.

Note: Error bars in the coefficient plots represent 95% confidence intervals.

its effects are disproportionately concentrated among female candidates across electoral contexts, with the effects being particularly pronounced at the local level.

Moreover, as shown in Figure 2, the gendered effects of online visibility are consistent across multiple forms of offensive speech. Overall, while the estimated average marginal effects of a one-hundred-thousand-follower increase are small and often negligible for men, they are substantially larger for women. In the parliamentary election, profanity increases only marginally for men (0.01 pp) but rises more noticeably for women (0.13 pp), and demeaning speech grows more than ten times as much for women (0.38 pp) as for men (0.04 pp). These gender gaps widen further in the local elections, where the estimated increase in profanity is 0.15 percentage points for women compared to 0.03 pp for men, and demeaning speech surges to 2.7 pp for women versus 0.27 pp for men. A similar pattern emerges for physical threats, which increase modestly for men but far more sharply for women – 0.01 pp versus 0.17 pp in the parliamentary election and 0.06 pp versus 0.39 pp in the local elections. In addition, across both elections, increases in ridicule-related offensive speech are modest for men but substantially larger for women – 0.02 pp versus 0.22 pp in the parliamentary election and 0.05 pp versus 1.47 pp in the local elections. Taken together, these findings indicate that online visibility disproportionately amplifies women’s exposure to all forms of offensive speech, with particularly pronounced effects in the local election context.

Figure 3 further shows that the gendered consequences of online visibility extend to identity-based hate speech across religious and ethnic dimensions, with particularly pronounced effects for ethnicity-based hate speech. In the parliamentary election, a one-hundred-thousand-follower increase corresponds to a 0.25 pp increase in religious hate speech targeting women, while the




Figure 3. Online visibility and gender differences across hate speech types.

Note: Error bars in the coefficient plots represent 95% confidence intervals.

estimated effect for men is negligible (0.0004 pp). In the local elections, the same increase is associated with a 0.63 pp rise for women, whereas the effect for men is small and negative (-0.007 pp). Ethnicity-based hate speech exhibits larger gender disparities. In the parliamentary election, women experienced a 0.11 pp increase compared to 0.008 pp for men, while in the local elections, women's exposure rose sharply by 1.64 pp, far exceeding the 0.25 pp increase observed for men. Taken together, these findings indicate that increased online visibility disproportionately exposes women to intersectional forms of hate.

Finally, we examine how party affiliation shapes gendered patterns of online violence. As shown in the first row of Figure 4, gender differences in predicted probabilities of online abuse during the parliamentary election vary across parties, with female candidates from the Green Left Party (YSP) standing out as facing particularly disproportionate levels of hostility. For offensive speech, YSP women (20.17%) rank among the most targeted, substantially exceeding the corresponding probabilities for women in the AKP (6.56%), CHP (10.44%), İYİ (11.40%), and MHP (7.25%). The pattern is similar in hate speech, where YSP women exhibit the highest predicted probability (7.34%, exceeding not only YSP men (5.93%) but also women in the AKP (0.88%), CHP (1.74%), İYİ (2.34%), and MHP (1.72%). These results underscore that while men tend to attract more abuse overall in several parties, YSP women are exceptional in receiving consistently higher levels of predicted offensive and hate speech compared to women in all other parties.

The local election analysis reinforces this pattern of heightened vulnerability among women from DEM (previously, YSP). In terms of offensive speech, DEM women (25.28%) are by far the most exposed, facing almost double the rate of DEM men (12.92%) and well above the levels predicted for women in other parties, such as AKP (3.42%), CHP (15.88%), İYİ (7.75%), and MHP (4.94%). Similarly, female candidates in DEM display a predicted probability of 6.66 percent for hate speech, a figure that not only almost doubles that of DEM men (3.79%) but also eclipses all other parties' female candidates, such as AKP (0.6%), CHP (1.14%), İYİ (0.98%), and MHP (0.89%). Thus, while the incidence of online violence differs across parties, the gap between DEM women and their counterparts – both across gender and across party lines – is consistently stark.

These findings indicate that female candidates affiliated with YSP/DEM face a particularly acute form of online hostility, reflecting the convergence of multiple reinforcing dynamics rather




Figure 4. Candidate partisanship and gender differences in online violence.

Note: Error bars in the coefficient plots represent 95% confidence intervals.

than any single source. The party's strong association with Kurdish identity and its prominent advocacy for feminist causes elevate the visibility of its female candidates beyond what party affiliation alone would predict. This heightened exposure appears to attract abuse that extends beyond generic sexist attacks to incorporate frames tied to political loyalty, security, and ethnic belonging—pointing to the salience of the Kurdish issue in structuring online hostility. Consistent with research showing that digital platforms reproduce entrenched gender stereotypes and intensify scrutiny of female politicians (Collignon and Rüdiger 2020; Håkansson 2021; Erikson, Håkansson, and Josefsson 2023), this visibility also activates norm-enforcing reactions that seek to delegitimize both their political authority and their public presence.

Across a series of robustness checks, our findings consistently indicate that candidate visibility has a stronger association with online violence directed at women than at men. First, adjusting the temporal scope of the analysis does not significantly alter our conclusions. As shown in Appendix 5, for the 2023 parliamentary election, we shorten the observation window from six months to three months and continue to observe statistically significant visibility effects for women relative to men. Moreover, Appendix 6 reports additional analyses restricted to candidates below the sample mean in tweet activity, with results broadly consistent with our main models, lending confidence that our findings are not driven by candidates whose high posting frequency may artificially inflate apparent visibility. Finally, as reported in Appendix 7, we replicate the regression analyses after excluding tweets directed at candidates at the upper and lower ends of the visibility distribution (top and bottom 5 percent based on follower counts), and the results remain similar to those of the main analysis.

Conclusion

This study shows that online visibility during election campaigns operates as a conditional political resource whose costs are distributed unevenly across male and female candidates. Drawing on nearly four million tweets directed at candidates in the 2023 parliamentary and 2024 local elections in Türkiye – an electoral autocracy where elections remain competitive but media access and political opportunity are systematically uneven – we demonstrate that increased online visibility is associated with disproportionately greater exposure to offensive and hate speech for female candidates. These gendered penalties appear across multiple forms of abuse, including ridicule, demeaning language, threats, and identity-based hate targeting ethnicity and religion. These patterns are especially pronounced among women affiliated with parties that occupy marginalized positions within the broader political landscape and explicitly promote gender equality.

These findings speak directly to one of the most consistent results in the existing literature: that visibility is a strong predictor of the abuse female politicians receive. Prior studies conducted in established democracies have repeatedly shown that women in more prominent political roles face disproportionately higher volumes of harassment and that this visibility penalty is markedly steeper for women than for men (Rheault, Rayment, and Musulan 2019; Gorrell, Bakir, Roberts et al. 2020; Håkansson 2021). Visibility in the existing literature has been operationalized through a range of proxies – list position, incumbency, political hierarchy, and media exposure (Gorrell, Bakir, Roberts et al. 2020; Håkansson 2021; Holm, Bjarnegård, and Zetterberg 2024) – each capturing a different dimension of political prominence. We contribute to this literature in several ways. First, by measuring visibility through follower counts, we provide a more precise test of the visibility-harassment relationship in the specific context where it matters most: digital platforms. Second, our findings corroborate the gendered visibility penalty across multiple forms of abuse – including ridicule, demeaning language, threats, and identity-based hate – thereby strengthening the robustness of this result with a new and more sophisticated measure. Finally, we extend this finding beyond consolidated democracies, showing it holds in an electoral autocracy where media constraints and patriarchal party structures make digital visibility both more necessary and more costly for female candidates.

By comparing national and local elections and incorporating partisan affiliation as a structural source of visibility, we also demonstrate that visibility penalties vary across electoral arenas and organizational contexts, with particularly strong effects in local elections where competition is more personalized. Our results show that female candidates affiliated with the YSP/DEM experience particularly intense online hostility due to the convergence of gendered, ethnic, and security-based stigmatization, reflecting both the securitization of Kurdish politics and backlash against women's visible leadership in a patriarchal political environment. Examining these dynamics in Türkiye highlights how regime and candidate characteristics, along with electoral context, shape the relationship between digital visibility and political vulnerability, suggesting that gendered harassment may intensify where electoral competition and public exposure coexist with weakened democratic safeguards.

Several limitations point to directions for future research. First, measuring online violence through handle-based mentions captures direct and public targeting of candidates but does not account for abuse occurring in private messages or through less visible practices such as coordinated harassment or disinformation campaigns. Second, while follower counts offer a transparent and readily observable indicator of candidates' public visibility, they capture only one dimension of a broader visibility construct. Other aspects of visibility – such as the extent to which candidates' messages are responded to, circulated, or otherwise engaged with, as well as the reach of campaign-related content – may also shape exposure to online harassment. Third, although this study leverages two elections within the same national context, comparative research across

countries and regime types would help clarify how broadly the visibility penalties identified here could extend beyond Türkiye. Longitudinal research, in addition, could shed light on how repeated exposure to online harassment shapes campaign strategies and female candidates' decisions to enter or exit politics.

Finally, our findings have implications for policies related to platform governance. Existing research demonstrates that platform design choices and moderation practices shape both the scale of such abuse and its political consequences (Citron 2014; Gillespie 2018). In the absence of effective intervention, online harassment can deter candidacies, constrain campaign behavior, and normalize intimidation as a political tactic (DiMeco 2021). The findings here point to the need for more robust and proactive measures to check abuse and for the consistent enforcement of safeguards against harassment, especially during election campaigns when visibility is most consequential. Strengthening these protections would help ensure that digital platforms do not reproduce gendered hierarchies and would move social media closer to realizing its potential as a more inclusive arena for women's political access, voice, and participation.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/S1475676526101650>.

Data availability statement. Replication files are available at <https://doi.org/10.7910/DVN/CPSFXD>.

Acknowledgements. We thank Anibal Perez-Liñan, Yaoyao Dai, participants of the 'Media Representation, Trolling, and Candidates' panel at the 2025 APSA Annual Meeting and the 'Strengthening Democratic Values in a Backsliding Era' symposium at the University of Pittsburgh, and the anonymous reviewers and editors at EJPR for their helpful comments on earlier drafts of this manuscript.

Author contributions. **Jungmin Han:** Conceptualization-Lead, Data curation-Lead, Investigation-Lead, Methodology-Equal, Project administration-Lead, Software-Equal, Supervision-Lead, Validation-Lead, Visualization-Lead, Writing - original draft-Lead, Writing - review & editing-Lead; **Müge Finkel:** Conceptualization-Equal, Funding acquisition-Lead, Project administration-Lead, Supervision-Equal, Writing - review & editing-Equal; **Steven Finkel:** Conceptualization-Equal, Investigation-Supporting, Methodology-Supporting, Supervision-Lead, Validation-Equal, Writing - review & editing-Lead; **Firat Duruşan:** Methodology-Equal, Software-Lead, Writing - original draft-Supporting, Writing - review & editing-Supporting; **Erdem Yörüük:** Methodology-Supporting, Resources-Supporting, Software-Supporting, Writing - review & editing-Supporting.

Financial support. This research received financial support from the Ford Institute for Human Security at the University of Pittsburgh. Jungmin Han acknowledges support from the European Research Council under the European Union's Horizon 2020 research and innovation program (No. 101002240). Müge Finkel acknowledges support through the Kabak Fellowship (2024-2027) from the Global Studies Center at the University of Pittsburgh.

Competing interests. The author(s) declare none.

AI use statement. The authors used AI tools only for basic grammar checks and sentence-level paraphrasing during the preparation of this manuscript.

References

- Alkumah, F. and Ma, X. (2022). 'A literature review of textual hate speech detection methods and datasets'. *Information*, 13(6), 273. <https://doi.org/10.3390/info13060273>.
- Alzahrani, M. and AlGhamdi F. (2025). 'Social media sentiment analysis for sustainable rural event planning: A case study of agricultural festivals in Al-Baha, Saudi Arabia'. *Sustainability*, 17(9), 3864. <https://doi.org/10.3390/su17093864>.
- Bardall, G. (2013). 'Gender-specific election violence: The role of information and communication technologies'. *Stability: International Journal of Security and Development*, 2(3), 60. <https://doi.org/10.5334/sta.cs>.
- Boulianne, S. (2020). 'Twenty years of digital media effects on civic and political participation'. *Communication Research*, 47(7), 947–966. <https://doi.org/10.1177/0093650218808186>.
- Cassese, E.C. and Barnes, T.D. (2019). 'Reconciling sexism and women's support for republican candidates: A look at gender, class, and whiteness in the 2012 and 2016 presidential races'. *Political Behavior*, 41(3), 677–700. <https://doi.org/10.1007/s11109-018-9468-2>.

- Cardo, V. (2021). 'Gender politics online? Political women and social media at election time in the United Kingdom, the United States and New Zealand.' *European Journal of Communication* 36(1), 38–52. <https://doi.org/10.1177/0267323120968962>.
- Chiluwa, I. (2022). 'Women's online advocacy campaigns for political participation in Nigeria and Ghana.' *Critical Discourse Studies*, 19(5), 465–484. <https://doi.org/10.1080/17405904.2021.1999287>.
- Cirimele, F., Macaluso, M., Agolino G., Çabuk Kaya N. and Zappulla C. (2025). 'Gender differences in the use of online platforms for political participation.' *Frontiers in Communication*, 10, article no. 1625965. <https://doi.org/10.3389/fcomm.2025.1625965>.
- Citron, D.K. (2014). *Hate Crimes in Cyberspace*. Cambridge, MA: Harvard University Press.
- Collignon, S. and Rüdiger, W. (2020). 'Harassment and intimidation of parliamentary candidates in the United Kingdom.' *The Political Quarterly*, 91(2), 422–429. <https://doi.org/10.1111/1467-923X.12855>.
- Delay, G. and Toremark, L. (2024). *Democracy in Turkey*. Chatham House, available at <https://www.chathamhouse.org/2024/05/democracy-turkey>, accessed 16 December 2024.
- Darı, A.B. (2018). 'Social media and politics: The use of social media by political parties in Turkey [Sosyal medya ve siyaset: Türkiye'deki siyasi partilerin sosyal medya kullanımı].' *Al Farabi International Journal of Social Sciences [Al Farabi Uluslararası Sosyal Bilimler Dergisi]*, 2(1), 1–10. <https://dergipark.org.tr/tr/pub/farabi/article/370664>.
- Diamond, L. (2002). 'Elections without democracy: Thinking about hybrid regimes.' *Journal of Democracy*, 13(2), 21–35.
- DiMeco, L. (2021). 'Gendered disinformation is a national security problem.' Brookings TechStream, March 8. Available at: <https://www.brookings.edu/techstream/>, accessed 20 June 2025.
- Dirlik, D. (2016). 'Erdogan's war on women.' *OpenDemocracy*, available at <https://www.opendemocracy.net/en/5050/erdogan-war-on-women/>, accessed 24 November 2024.
- Eagly, A.H. and Karau, S.J. (2002). 'Role congruity theory of prejudice toward female leaders.' *Psychological Review*, 109(3), 573–598. <https://doi.org/10.1037/0033-295X.109.3.573>.
- Erikson, J., Håkansson, S. and Josefsson, C. (2023). 'Three dimensions of gendered online abuse: Analyzing Swedish MPs' experiences of social media.' *Perspectives on Politics*, 21(3), 896–912. <https://doi.org/10.1017/S1537592721002048>.
- Erzurum, F. (2013). 'Representation of women in Turkish television's primetime news.' *Online Journal of Communication and Media Technologies*, 3(2), 162–186. <https://doi.org/10.29333/ojcm/2428>.
- Finkel, M., Thakur, D., Finkel, S., Duruşan, F., Yörük, E., Topçu, I., Yardi, M.C., Zaner, A. and Han, J. (2025). 'Intersectional hatred: An application of large language models to detect hate and offensive speech targeted at congressional candidates in the 2024 US election.' *Companion Proceedings of the ACM on Web Conference 2025*, 2776–2781. <https://doi.org/10.1145/3701716.3716880>.
- Fortuna, P. and Nunes, S. (2018). 'A survey on automatic detection of hate speech in text.' *ACM Computing Surveys*, 51(4), 1–30. <https://doi.org/10.1145/3232676>.
- Gillespie, T. (2018). *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press.
- Gorrell, G., Bakir, M.E., Roberts, I., Greenwood, M.A. and Bontcheva, K. (2020). 'Online abuse toward candidates during the UK General Election 2019.' arXiv preprint, arXiv:2001.08686. <https://doi.org/10.48550/arXiv.2001.08686>.
- Gunitsky, S. (2015). 'Corrupting the cyber-commons: Social media as a tool of autocratic stability.' *Perspectives on Politics*, 13(1), 42–54. <https://doi.org/10.1017/S1537592714003120>.
- Håkansson, S. (2021). 'Do women pay a higher price for power? Gender bias in political violence in Sweden.' *The Journal of Politics*, 83(2), 515–531. <https://doi.org/10.1086/709838>.
- Hawkesworth, M. (2003). 'Congressional enactments of race-gender: Toward a theory of raced-gendered institutions.' *American Political Science Review*, 97(4), 529–550. <https://doi.org/10.1017/S0003055403000868>.
- Hickey, D., Schmitz, M., Fessler, D., Smaldino, P.E., Muric, G. and Burghardt, K. (2023). 'Auditing Elon Musk's impact on hate speech and bots.' *Proceedings of the International AAAI Conference on Web and Social Media*, 17, 1133–1137. <https://doi.org/10.1609/icwsm.v17i1.22222>.
- Holm, M., Bjarnegård, E. and Zetterberg, P. (2024). 'Comparing gendered exposure and impact in online election violence: Tunisian political candidates targeted on Facebook.' *Politics & Gender*, 20(4), 879–902. <https://doi.org/10.1017/S1743923X24000163>.
- Hürriyetoglu, A., Tanev, H., Zavarella, V., Piskorski, J., Yeniterzi, R., Mutlu, O., Yuret, D. and Villavicencio, A. (2021). 'Challenges and applications of automated extraction of socio-political events from text (CASE 2021): Workshop and shared task report.' *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, 1–9. <http://doi.org/10.18653/v1/2021.case-1.1>.
- İkiz, O.O., Sobaci, M.Z., Yavuz, N. and Karkin, N. (2014). 'Political use of Twitter: The case of metropolitan mayor candidates in 2014 local elections in Turkey.' *Proceedings of the 8th International Conference on Theory and Practice of Electronic Governance*, 41–50. <https://doi.org/10.1145/2691195.2691219>.
- International Foundation for Electoral Systems (IFES). (2018). *Women and Technology: Enhancing Political Participation*. IFES, available at <https://www.ifes.org/news/women-and-technology-enhancing-political-participation>, accessed 24 June 2025.

- Inter-Parliamentary Union (IPU). (2024). *Women in Parliament 2023*. Inter-Parliamentary Union, available at <https://www.ipu.org/resources/publications/reports/2024-03/women-in-parliament-2023>, accessed 24 June 2025.
- Junn, J. and Masuoka, N. (2020). 'The gender gap is a race gap: Women voters in US presidential elections.' *Perspectives on Politics*, 18(4), 1135–1145. <https://doi.org/10.1017/S1537592719003876>.
- Kemp, S. (2025). "Digital 2025: Turkey." DataReportal. March 3. available at <https://datareportal.com/reports/digital-2025-turkey>, accessed 15 June 2025.
- Krook, M.L. (2017). 'Violence against women in politics.' *Journal of Democracy*, 28(1), 74–88. <https://doi.org/10.1353/jod.2017.0007>.
- Krook, M.L. (2020). 'The cost of doing politics? Analyzing violence and harassment against female politicians.' *Perspectives on Politics*, 18(3), 740–755. <https://doi.org/10.1017/S1537592719001397>.
- Krook, M.L. and Restrepo Sanín, J. (2016). 'Gender and political violence in Latin America: Concepts, debates and solutions.' *Política y Gobierno*, 23(1), 127–162.
- Matsilele, T. and Nkoala, S. (2023). 'Metavoicing, trust-building mechanisms and partisan messaging: A study of social media usage by selected South African female politicians.' *Information, Communication & Society*, 26(13), 2575–2597. <https://doi.org/10.1080/1369118X.2023.2252862>.
- Newman, N., Fletcher, R., Robertson, C.T., Ross Arguedas, A. and Nielsen, R.K. (2024). *Reuters Institute Digital News Report 2024*. Reuters Institute for the Study of Journalism. available at <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2024>, accessed 1 May 2025.
- OSCE Office for Democratic Institutions and Human Rights (ODIHR). (2023). 'Republic of Türkiye: General Elections 14 May and Presidential Election, Second Round 28 May 2023.' ODIHR Election Observation Mission Final Report. Warsaw: OSCE/ODIHR, available at <https://odihr.osce.org/odihr/elections/turkiye/537642>, accessed 23 June 2025.
- Ortiz, R., Ahmed, S., Salas, P. and Eisty, N.U. (2025). 'Analyzing social media engagement of computer science conferences.' arXiv preprint, arXiv:2503.01038. <https://doi.org/10.48550/arXiv.2503.01038>.
- Osei-Appiah, S., Mutsvairo, B. and Orgeret, K.S. (2023). 'Women and digital political communication in non-Western societies.' *Information, Communication & Society*, 26(13), 2507–2517. <https://doi.org/10.1080/1369118X.2023.2257995>.
- Pacilli, M.G. and Mannarini, T. (2019). 'Are women welcome on Facebook? A study of Facebook profiles of Italian female and male public figures.' *TPM: Testing, Psychometrics, Methodology in Applied Psychology*, 26(2), 199–207. <https://doi.org/10.4473/TPM26.2.2>.
- Phillips, A. (1995). *The Politics of Presence*. Clarendon Press.
- Puwar, N. (2004). *Space Invaders: Race, Gender and Bodies out of Place*. Berg.
- Reuters. (2024). "Turkey opposition aims to hit back at Erdogan in local elections." Reuters, March 28, available at <https://www.reuters.com/world/middle-east/turkey-opposition-aims-hit-back-erdogan-local-elections-2024-03-28/>, accessed 24 May 2025.
- Rheault, L., Rayment, E. and Musulan, A. (2019). 'Politicians in the line of fire: Incivility and the treatment of women on social media.' *Research & Politics*, 6(1), article no. 2053168018816228. <https://doi.org/10.1177/2053168018816228>.
- Richardson-Self, L. (2021). *Hate Speech Against Women Online: Concepts and Countermeasures*. Bloomsbury Publishing.
- Rudman, L.A., Moss-Racusin, C.A., Phelan, J.E. and Nauts, S. (2012). 'Status incongruity and backlash effects: Defending the gender hierarchy motivates prejudice against female leaders.' *Journal of Experimental Social Psychology*, 48(1), 165–179. <https://doi.org/10.1016/j.jesp.2011.10.008>.
- Schedler, A. (2002). 'Elections without democracy: The menu of manipulation.' *Journal of Democracy*, 13(2), 36–50. <https://doi.org/10.1353/jod.2002.0031>.
- Silverman, R. (2024). 'Local elections were a loss for the AKP, were they a victory for the opposition?' *The Türkiye Analyst*, available at <https://www.turkeyanalyst.org/publications/turkey-analyst-articles/item/721-local-elections-were-a-loss-for-the-akp-were-they-a-victory-for-the-opposition?.html>, accessed 24 July 2025.
- Sinpeng, A. (2021). 'Hashtag activism: Social media and the #FreeYouth protests in Thailand.' *Critical Asian Studies*, 53(2), 192–205. <https://doi.org/10.1080/14672715.2021.1882866>.
- Sobaci, M.Z. and Karkin, N. (2013). 'The use of Twitter by mayors in Turkey: Tweets for better public services?' *Government Information Quarterly*, 30(4), 417–425. <https://doi.org/10.1016/j.giq.2013.10.003>.
- Southern, R. and Harmer, E. (2021). 'Twitter, incivility and "everyday" gendered othering: An analysis of tweets sent to UK Members of Parliament.' *Social Science Computer Review*, 39(2), 259–275. <https://doi.org/10.1177/0894439319865519>.
- Theocharis, Y., Barberá, P., Fazekas, Z., Popa, S.A. and Parnet, O. (2016). 'A bad workman blames his tweets: The consequences of citizens' uncivil Twitter use when interacting with party candidates.' *Journal of Communication*, 66(6), 1007–1031. <https://doi.org/10.1111/jcom.12259>.
- Ugur-Cinar, M. (2023). 'Elections and democracy in Turkey: Reconsidering competitive authoritarianism in the age of democratic backsliding.' *The Political Quarterly*, 94(3), 445–451. <https://doi.org/10.1111/1467-923X.13309>.

- United Nations.** (2020). *Guidance on Addressing Hate Speech in the Field: Implementation of the United Nations Strategy and Plan of Action on Hate Speech*. United Nations, available at https://digitallibrary.un.org/record/3889286/files/UN_Strategy_and_PoA_on_Hate_Speech_Guidance_on_Addressing_in_field.pdf, accessed 24 May 2024.
- Vaccari, C. and Valeriani, A.** (2021). *Outside the Bubble: Social Media and Political Participation in Western Democracies*. Oxford University Press.
- Yarchi, M. and Samuel-Azran, T.** (2018). 'Women politicians are more engaging: Male versus female politicians' ability to generate users' engagement on social media during an election campaign.' *Information, Communication & Society*, **21**(7), 978–995.